Predictive Simulations of Particle-laden Turbulence in a Radiation Environment

Gianluca Iaccarino
Mechanical Engineering Department
Institute for Computational Mathematical Engineering
Stanford University

Stanford
PSAAP II

a collaboration between

Stanford University - University of Michigan - SUNY Stony Brook
University of Minnesota - University of Colorado - University of Texas

Stanford PSAAP II

Mission:

To develop and demonstrate predictive multi-physics simulations of particle-laden turbulence subject to radiation on next-generation exascale compute systems

Predictive Science Academic Alliance Program

Stanford
PSAAP II

Mission:

To develop and demonstrate predictive multi-physics simulations of particle-laden turbulence subject to radiation on next-generation exascale compute systems

Predictive Science Academic Alliance Program

Relevant Applications

Solar
Thermal collectors

Soot generation in
engines

Solid-fuel
rocket
exhausts

Accidental fires

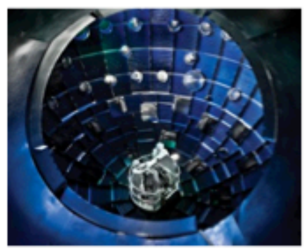

Inertial confinement
fusion

Atmospheric entry
thermal protection systems

Multiphysics Scenario

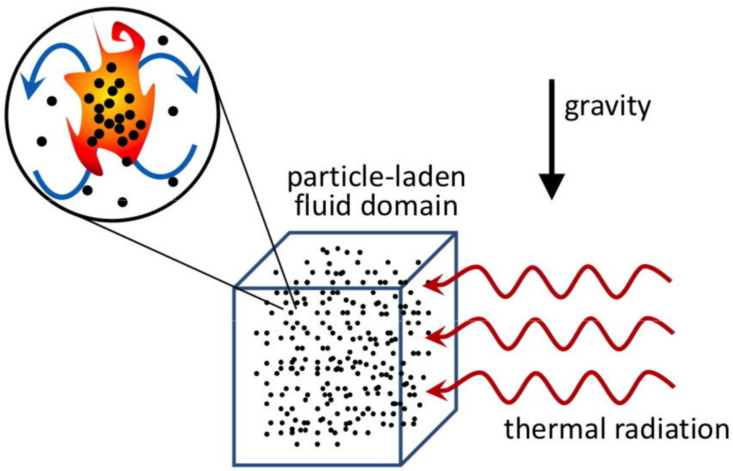

In the absence of external forcing turbulence in particle-laden flows decays…what happens with radiation ?

Multiphysics Scenario

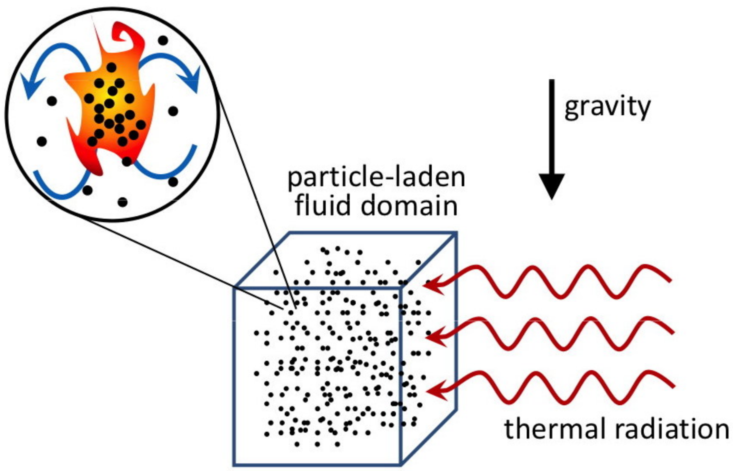

In the absence of external forcing turbulence in particle-laden flows decays…what happens with radiation ?

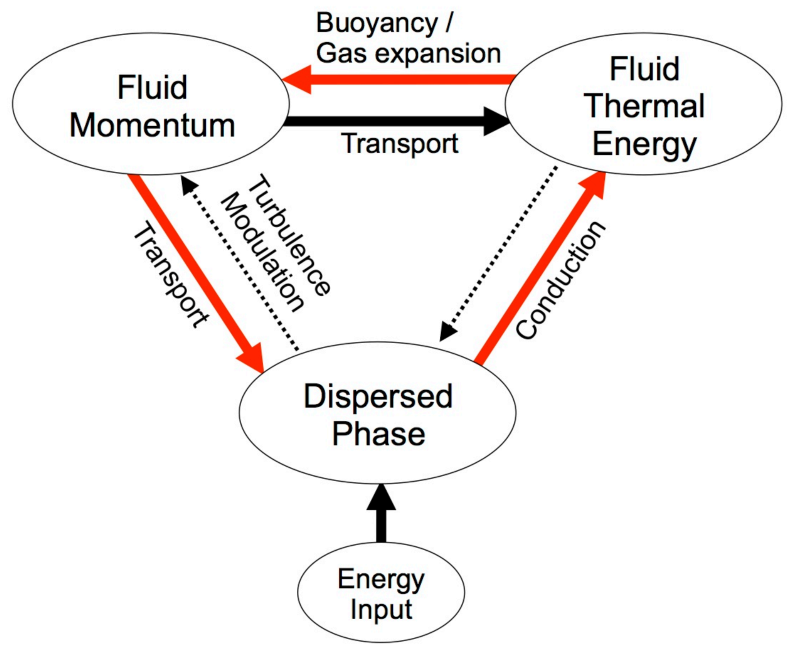

The Center focus is on the three-way coupled regime

Scientific Challenge

Multi-physics coupling of: turbulence, radiation, and particle transport

Starting point is particle-laden turbulence

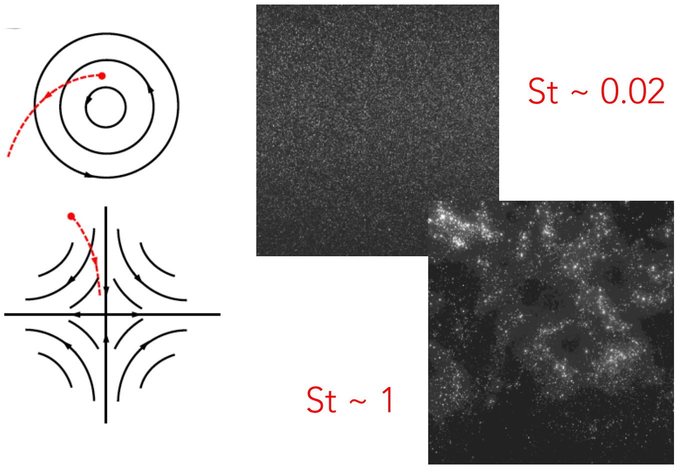

Clustering (preferential concentration) in particle-laden turbulence - J. Eaton
Turbulence Modulation

Scientific Challenge

Multi-physics coupling of: turbulence, radiation, and particle transport

Starting point is particle-laden turbulence

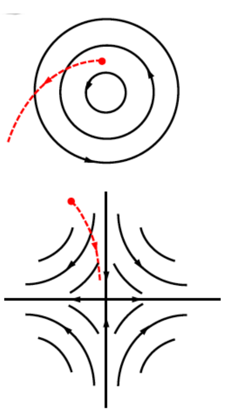

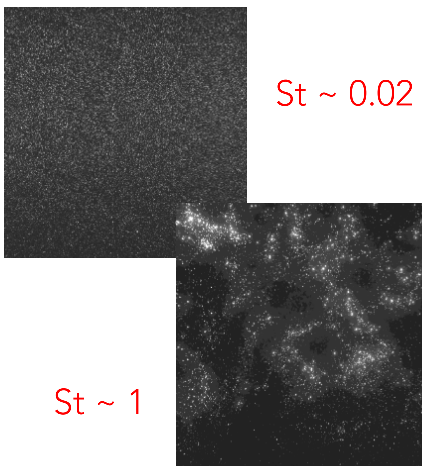

St ~ 0.02

St ~ 1

Clustering (preferential concentration) in particle-laden turbulence - J. Eaton

Turbulence Modulation

Add Radiation

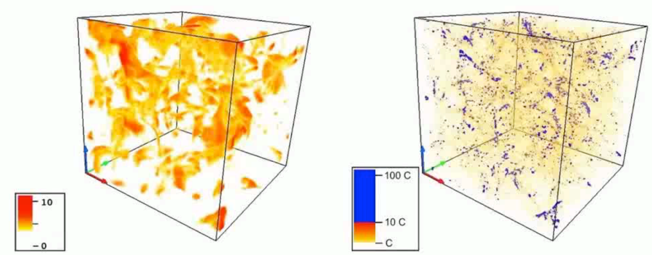

Temperature

St ~ 1

Particle Concentrations

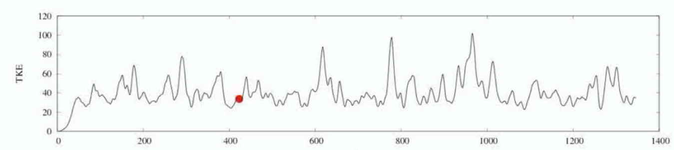

A-ha Moment #1: Turbulence is sustained!

R. Zamanski, H. Poraunsari

What are we trying to predict?

- Global validation quantities
  - –Radiation power transmission
  - –Integral gas energy gain
  - –Exit mean temperature profile
  - –Temperature fluctuation intensity
- Local validation quantities
  - –Mean velocity profile for gas/particles
  - –Instantaneous particle concentration (2D slices)
  - –Gas turbulence properties

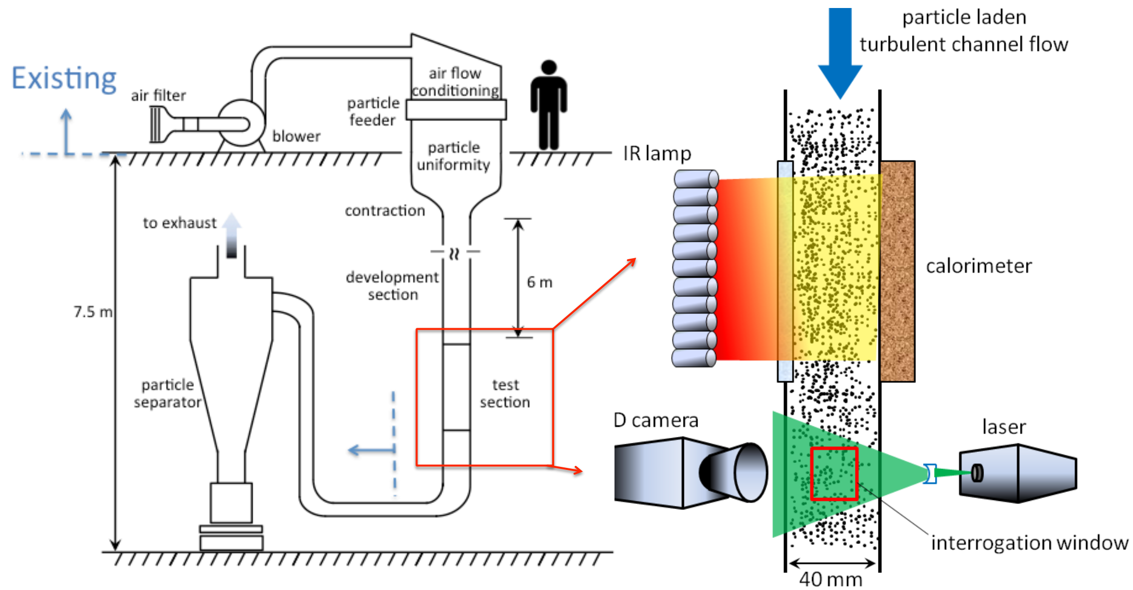

Eaton vertical tunnel @ Stanford

Research Agenda

Single-Physics Modeling

- Investigate radiation models
  - (in)homogeneous absorption
  - discrete ordinates vs. ray tracing
- Beyond point-particle tracking
  - finite-size effects
  - wall interactions, collisions

Coupling effects

- Radiation absorption by particles
  - Emissivity, surface properties, secondary absorption
- Heat transfer between particles & fluid
  - non-equilibrium, non-locality
- Turbulence
  - Modulaton by particles: dissipation in wakes, …
- Three-way interactions

Research Agenda

Single-Physics Modeling

- Investigate radiation models
  - (in)homogeneous absorption
  - discrete ordinates vs. ray tracing
- Beyond point-particle tracking
  - finite-size effects
  - wall interactions, collisions

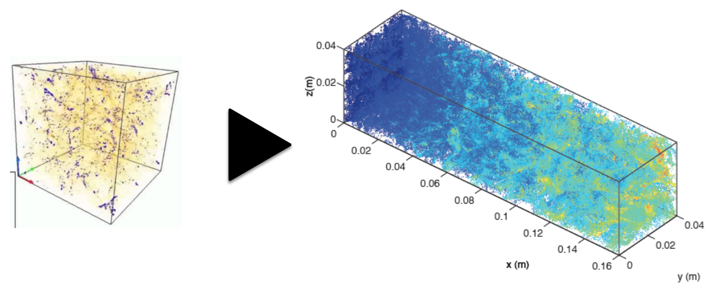

Coupling effects

- Radiation absorption by particles
  - Emissivity, surface properties, secondary absorption
- Heat transfer between particles & fluid
  - non-equilibrium, non-locality
- Turbulence
  - Modulaton by particles: dissipation in wakes, …
- Three-way interactions

Present Simulations

Turbophoresis:

- particles migrate in the direction of decreasing turbulence
- Increased residence time near walls >> enhanced radiation absorption >> higher temperatures

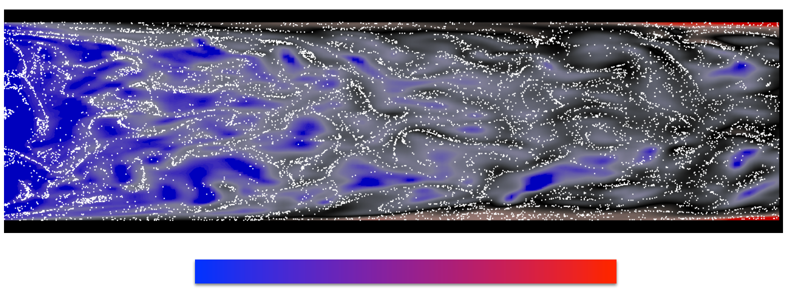

A-ha Moment #2: Wall interactions are of primary importance for experimental apparatus design!

Deposition:

- turbulence-driven secondary motions in corners
- preferential concentration & preferential deposition

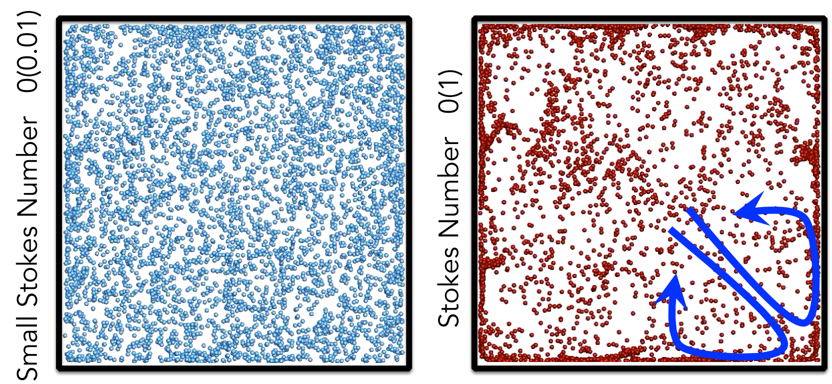

M. Esmaily, H. **Abdehkakha**

Uncertainty Quantification

Many sources of uncertainties

Naturally occurring

- Particle size/property variability
- Radiation forcing
- Losses through walls
- Inflow/Injection conditions

…

Mathematical models

- Particle physics
- Radiation/particle coupling
- Airflow/particle interactions
- …

Uncertainty Quantification

Many sources of uncertainties

Naturally occurring

- Particle size/property variability
- Radiation forcing
- Losses through walls
- Inflow/Injection conditions

…

Mathematical models

- Particle physics
- Radiation/particle coupling
- Airflow/particle interactions
- …

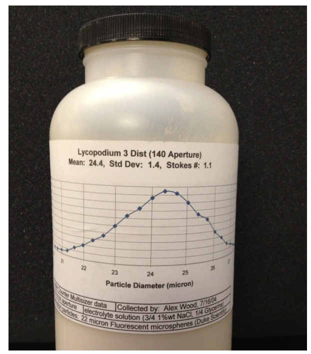

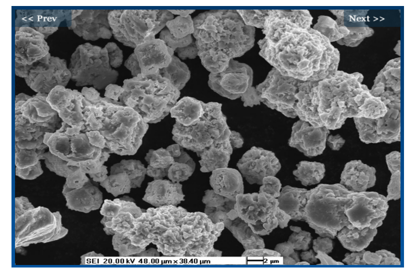

Particle Size Distribution (Microtrac)with Sonication	
% Passing	Microns
D10	4.62
D20	5.71
D30	6.60
D40	7.45
D50	8.33
D60	9.31
D70	10.47
D80	12.07
D90	14.81

A microscopic view of the nickel powder used in PSAAP II

Particles consist of non-spherical shapes with variation in size, indicating strong polydispersity.

Uncertainty Quantification

Many sources of uncertainties

Naturally occurring

- Particle size/property variability
- Radiation forcing
- Losses through walls
- Inflow/Injection conditions

…

Mathematical models

- Particle physics
- Radiation/particle coupling
- Airflow/particle interactions
- …

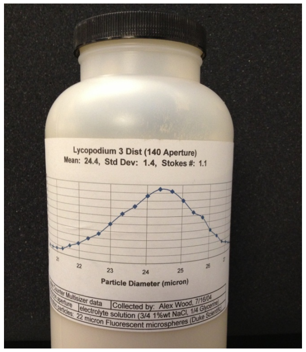

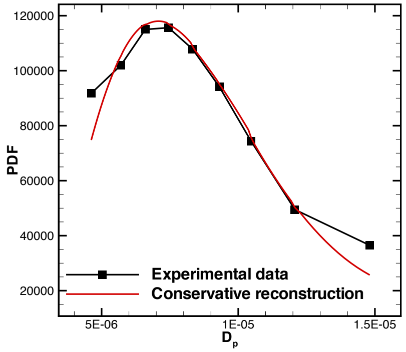

UQ-perspective on polydispersity

Step 1: Construct probability density functions (PDFs) of particles size base on available experimental data.

Uncertainty Quantification

DNS Simulations

Mono-disperse

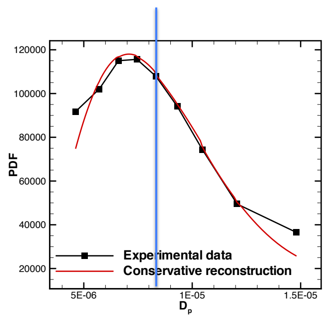

Poly-disperse

particle temperature

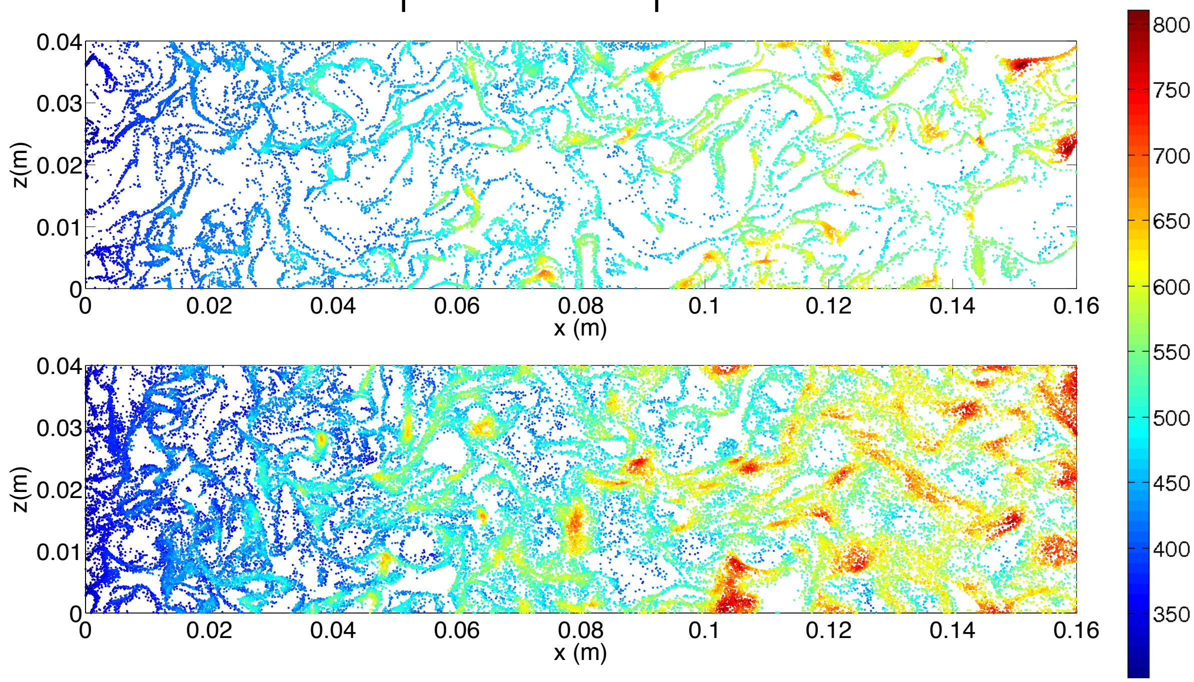

Uncertainty Quantification

DNS Simulations

Mono-disperse

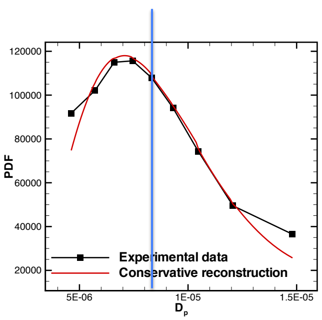

Poly-disperse

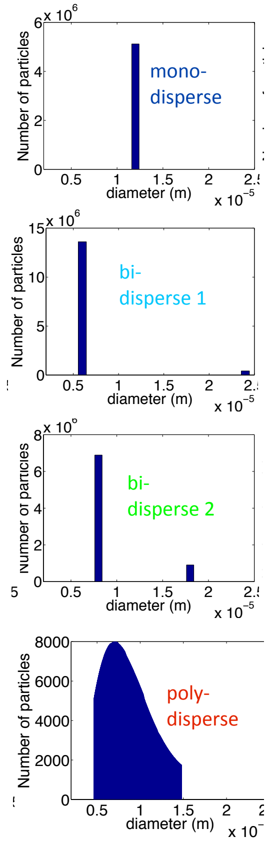

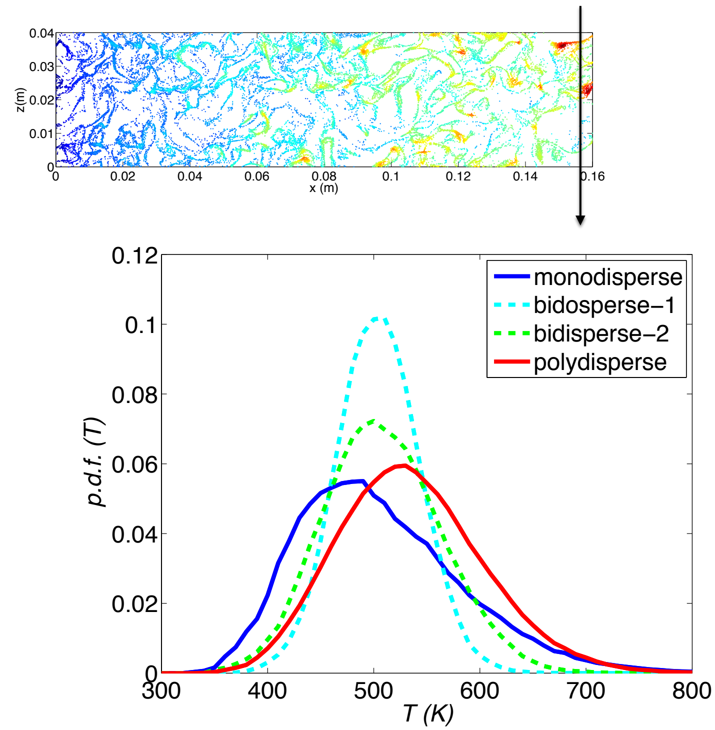

Gas temperature at the outlet
(note change in the mean T_f)

Uncertainty Quantification

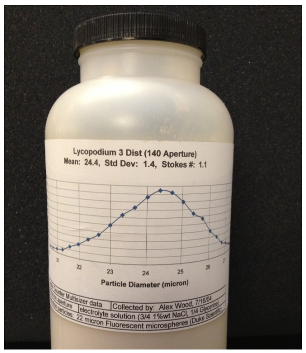

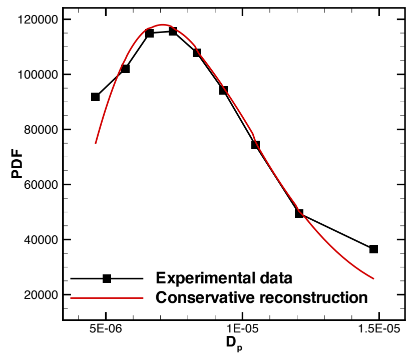

Uncertainty Quantification

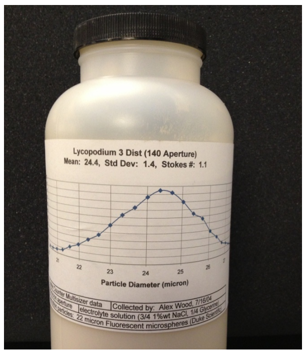

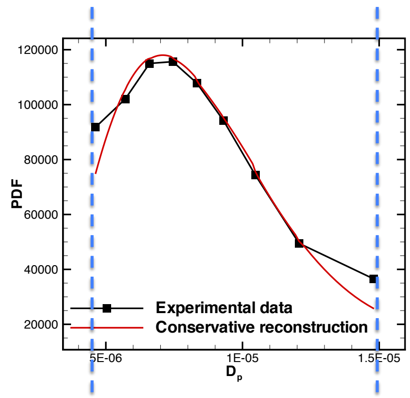

Experimental PDFs are truncated at small and large particle sizes resulting in uncertainty

Uncertainty Quantification

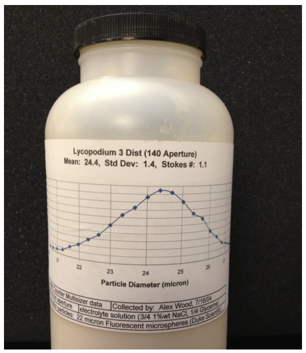

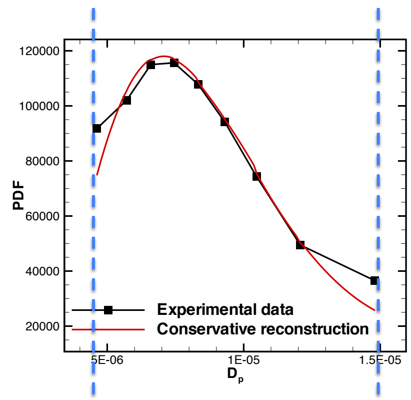

Experimental PDFs are truncated at small and large particle sizes resulting in uncertainty

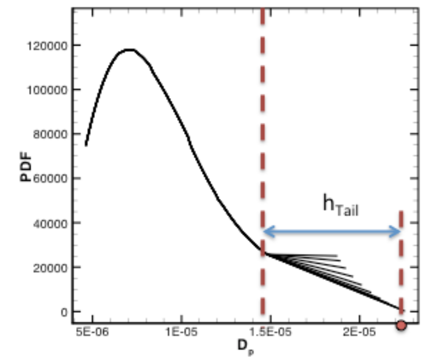

We “model” the tail as an additional uncertainty

All the cases are at the same mass loading, i.e. adding the tail redistributes the number of particles among the classes

Simulations based on 1D ROM

Uncertainty Quantification

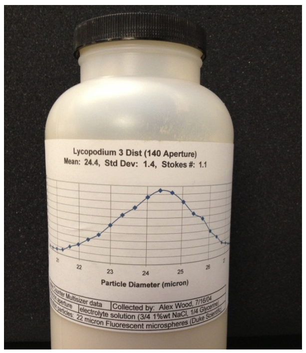

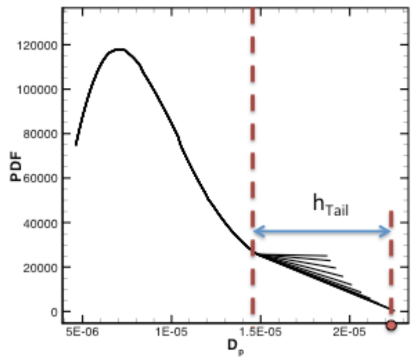

We "model" the tail as an additional uncertainty

All the cases are at the same mass loading, i.e. adding the tail redistributes the number of particles among the classes

Simulations based on 1D ROM

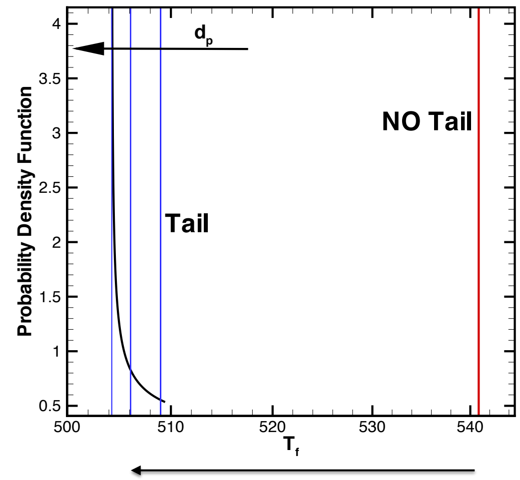

Effect of the tail uncertainty

A-ha moment #3: large particles reduce the overall absorption…

M. Rahmani, G. Geraci

Stanford PSAAP II

Mission:

To develop and demonstrate predictive multi-physics simulations of particle-laden turbulence subject to radiation on next-generation exascale compute systems

Predictive Science Academic Alliance Program

Stanford
PSAAP II

Mission:

To develop and demonstrate predictive multi-physics simulations of particle-laden turbulence subject to radiation on next-generation exascale compute systems

Predictive Science Academic Alliance Program

HPC: The Exascale Challenge

Complexity deriving from:

- Hybrid/Heterogeneous Nodes (CPUs, GPUs…)
- Deep Memory Hierarchies
- On-Chip vs In-Cabinet vs Across-System Communications
- Custom Non-conventional Networks
- Massive data management & mining
- Hardware Failures
- …

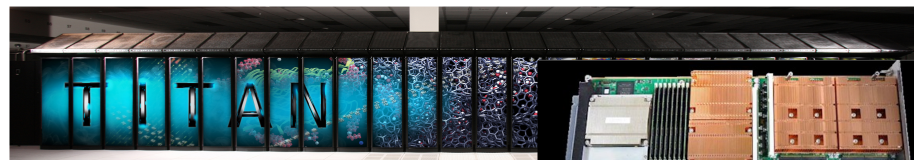

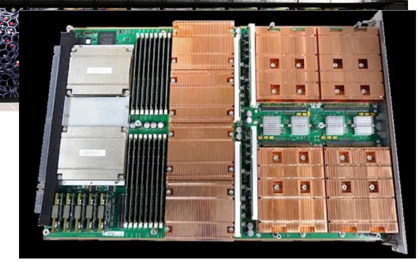

Titan

18,688 nodes (16-core AMD CPU + Nvidia K20X GPU)

Programming Barriers:

- Languages: CPU/GPU: C++, CUDA, OpenCL
- Queues for scheduling work between CPU/GPU?
- Nodes: Threads and locks – pthreads, OpenMP,…
- Cluster: Message passing – MPI
- Data management
- I/O
- Resiliency
- …

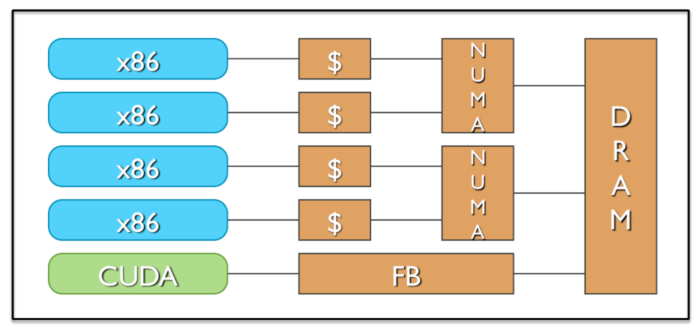

The road to exascale…

Our multiphysics problem naturally involves diverse computational strategies

- PDE solvers (flow)
- Lagrangian tracking (particles)
- Integro-differential equation (radiation)
- Intrusive and non-intrusive UQ approaches

and several options

- Eulerian transport for particles
- Ray tracing for radiation
- ….

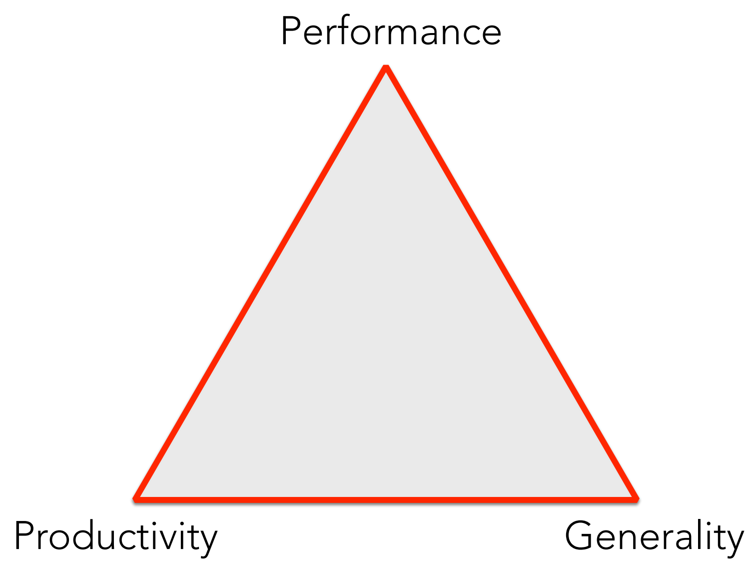

The road to exascale…

Our multiphysics problem naturally involves diverse computational strategies

- PDE solvers (flow)
- Lagrangian tracking (particles)
- Integro-differential equation (radiation)
- Intrusive and non-intrusive UQ approaches

and several options

- Eulerian transport for particles
- Ray tracing for radiation
- ….

Stanford PSAAP Program

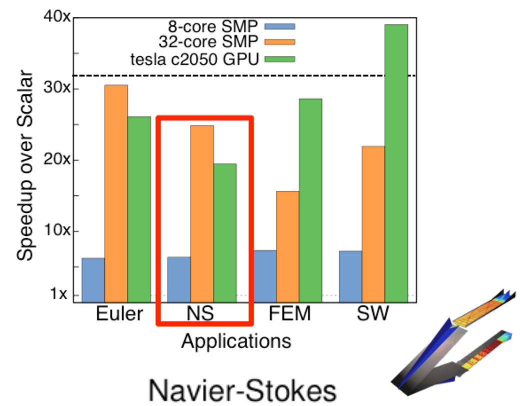

Navier-Stokes

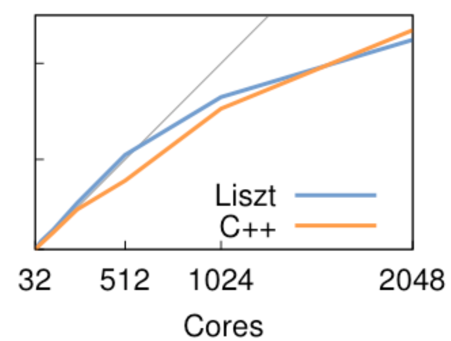

Since PSAAP…

Lessons Learned

- Interoperate, Interoperate, Interoperate
  - Call legacy libraries: e.g. system solvers
  - Work incrementally: drop DSL code into apps
  - A generalized DSL is not a DSL: multiple interacting DSLs
- DSL need to be easy and flexible
  - Embed DSL into existing languages: e.g. library rather than language
  - Dynamic meshes and adaptivity
  - Auto-tuning, JIT

Leverage concurrent co-design efforts (Hanrahan & Aiken)

- ExMatEx
  - Design and implementation of Terra/Lua
- ExaCT
  - Design and implementation of Legion

Key ingredients of PSAAP2

- Interoperable Domain Specific Languages
- Legion Automatic Memory Mapping/Management
- Resiliency Libraries

Since PSAAP…

Lessons Learned

- Interoperate, Interoperate, Interoperate
  - Call legacy libraries: e.g. system solvers
  - Work incrementally: drop DSL code into apps
  - A generalized DSL is not a DSL: multiple interacting DSLs
- DSLs need to be easy and flexible
  - Embed DSL into existing languages: e.g. library rather than language
  - Dynamic meshes and adaptivity
  - Auto-tuning, JIT

Leverage concurrent co-design efforts (Hanrahan & Aiken)

- ExMatEx
  - Design and implementation of Terra/Lua
- ExaCT
  - Design and implementation of Legion

Key ingredients of PSAAP2

- Interoperable Domain Specific Languages
- Legion Tasks & Memory Mapping/Management
- Resiliency Libraries

PSAAP II

Liszt >> Liszt 2.0
Scala >> Terra/Lua + Legion
>> CD

Liszt 2.0

Flexible

- Structured & unstructured grids + particles >> relations

Interoperability

- Lua/Terra Liszt integration complete
- Runtimes for CPU/GPU developed
- Lulesh test achieved 10x speedup (CPU2GPU) with no tuning

Liszt 2.0

Flexible
- Structured & unstructured grids + particles >> relations

Interoperability
- Lua/Terra Liszt integration complete
- Runtimes for CPU/GPU developed
- Lulesh test achieved 10x speedup (CPU2GPU) with no tuning

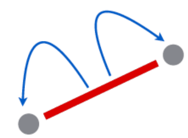

C/Fortran

```
int e_head[n_edges];
int e_tail[n_edges];
double v_temperature[n_verts];
```

The programmer explicitly defines the connection between vertices & edges and the compiler only "sees" integer pointers

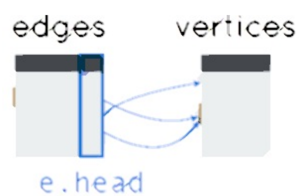

Liszt

```
liszt kernel(e : M.edges)
 var v1 = e.head
 var v2 = e.tail
```

The language implicitly provides relations and the compiler can interpret to determine data stencils

The Framework: a graphical view

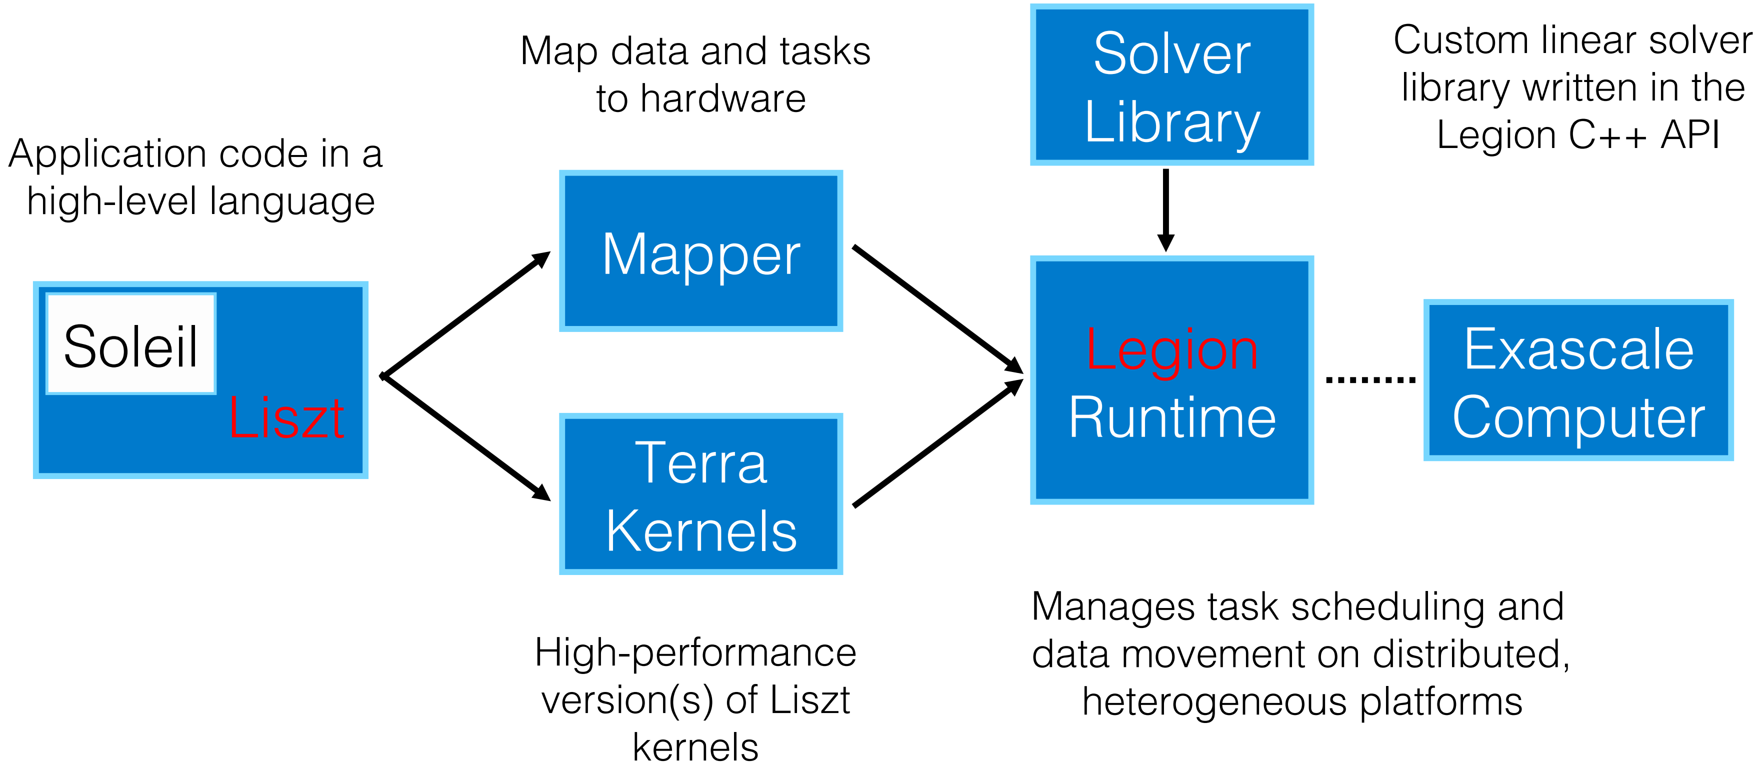

The Framework: a graphical view

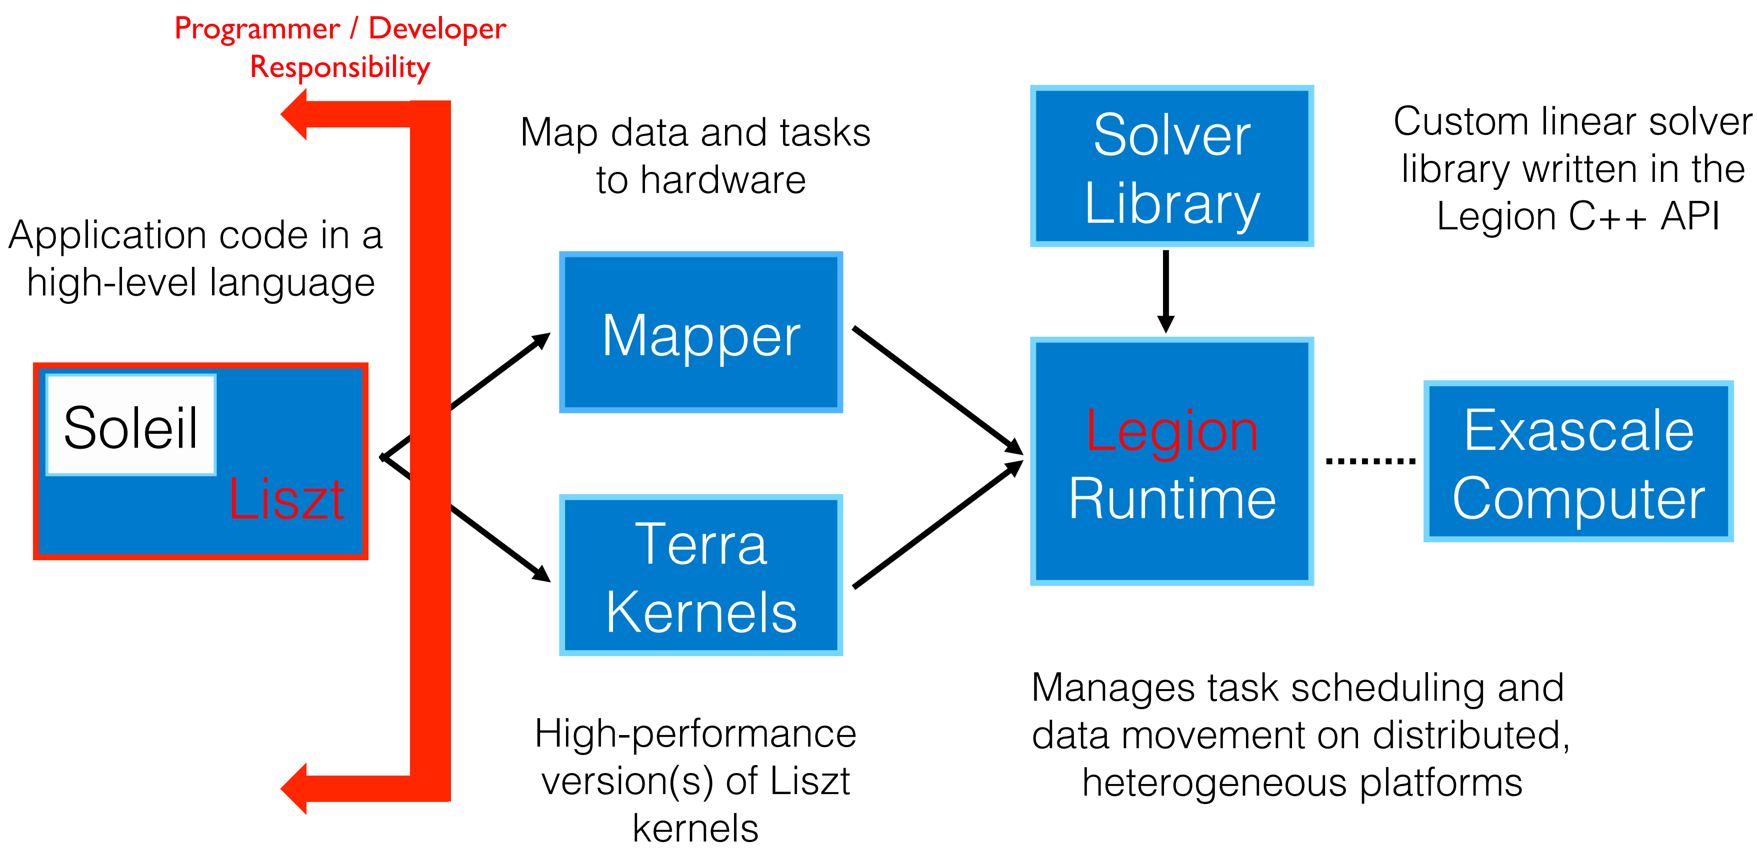

Soleil Development Timeline

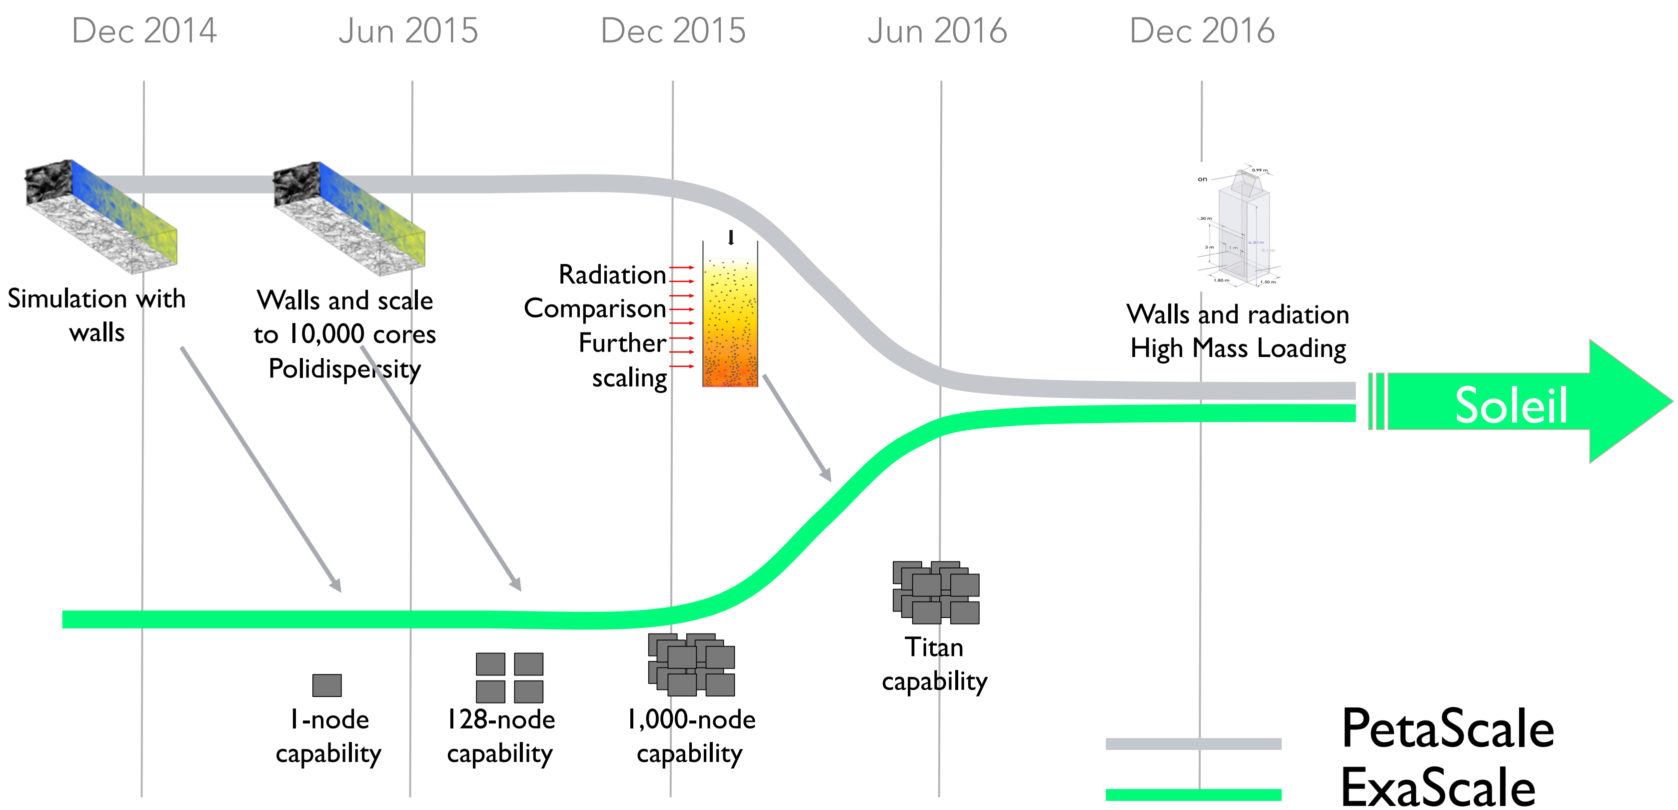

Soleil Today….

Soleil

- Completely written in Liszt 2.0
- 3D, Turbulence + Particles + Radiation

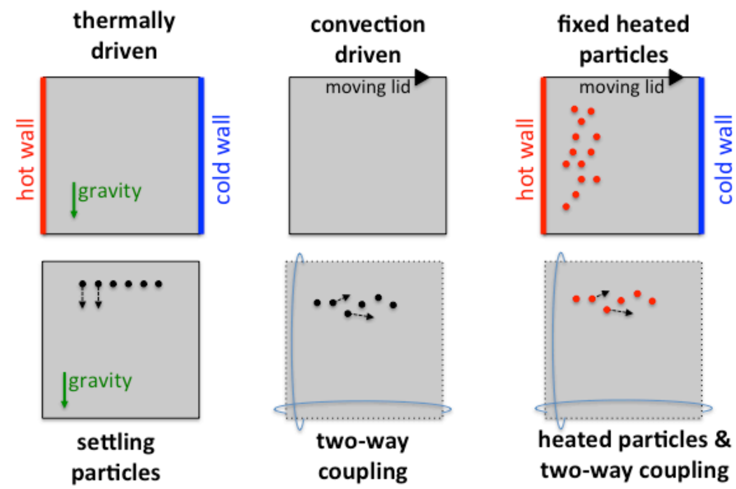

Verification & Regression Test Suite

Soleil Today….

Soleil

- Completely written in Liszt 2.0
- 3D, Turbulence + Particles + Radiation

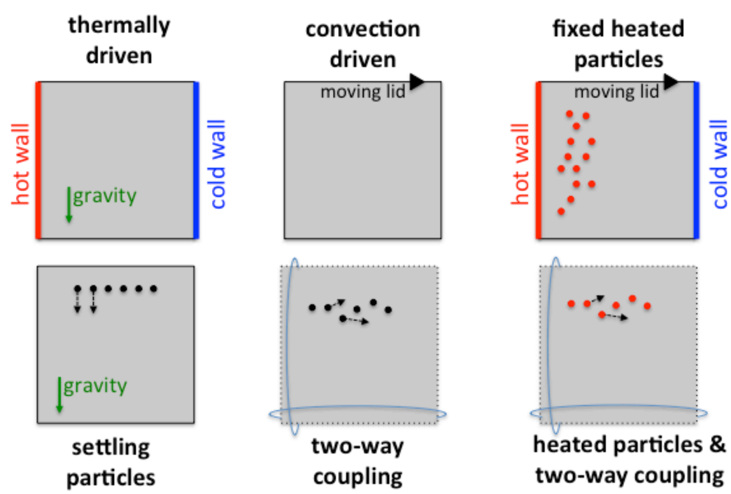

Verification & Regression Test Suite

Switching to GPU as simple as adding:

```
L.default_processor = L.GPU
```

Or switch to GPU for parts of an application:

```
relation:MoveTo(L.GPU)
L.default_processor = L.GPU
```

A-ha moment #4: Soleil test achieved 19x speedup (CPU2GPU) with no tuning

G. Bernstein, C. Lemire, T. Economon

The Framework: a graphical view

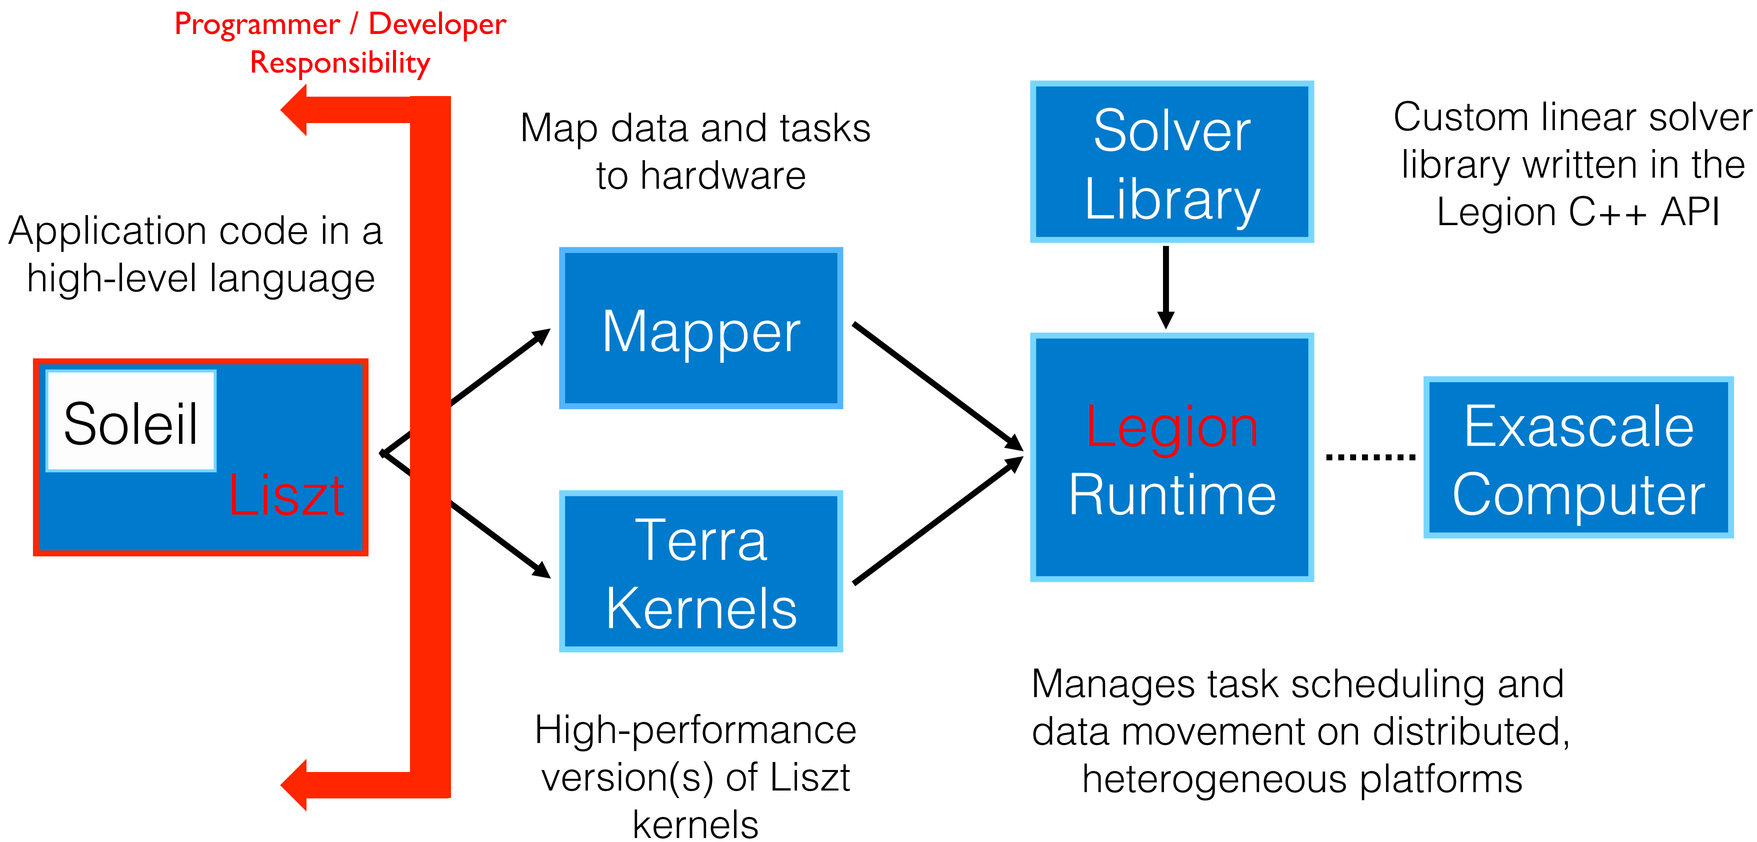

The Framework: a graphical view

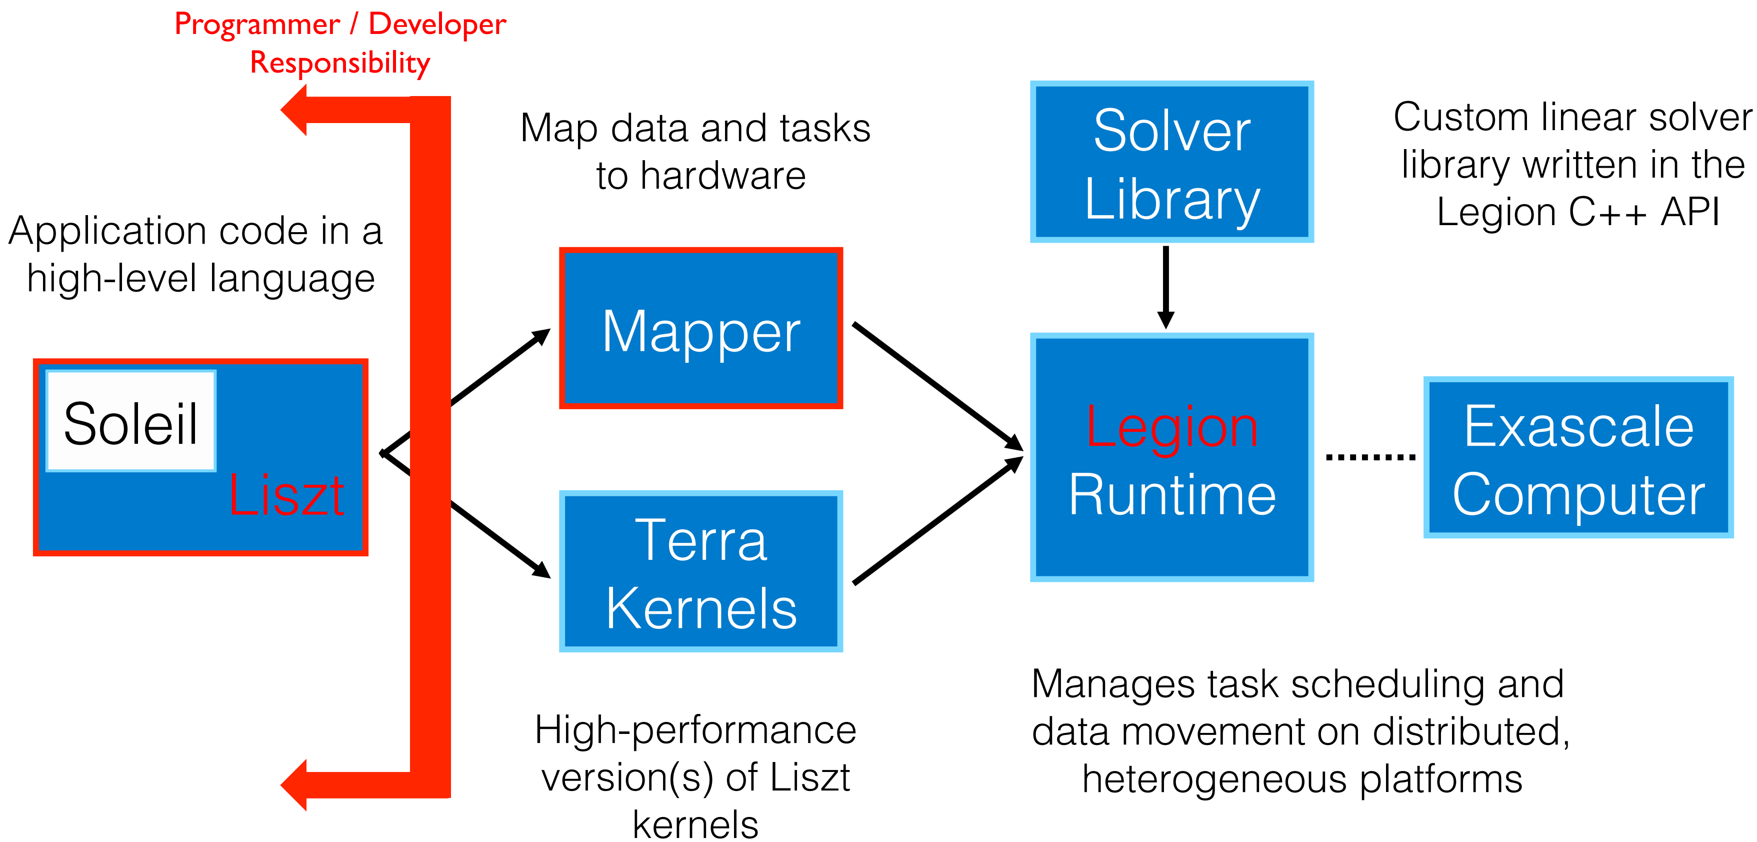

Legion Mapping Interface

Mapper/Programmer selects:
  Where tasks run
  Where regions (data) are placed

Mapping computed dynamically

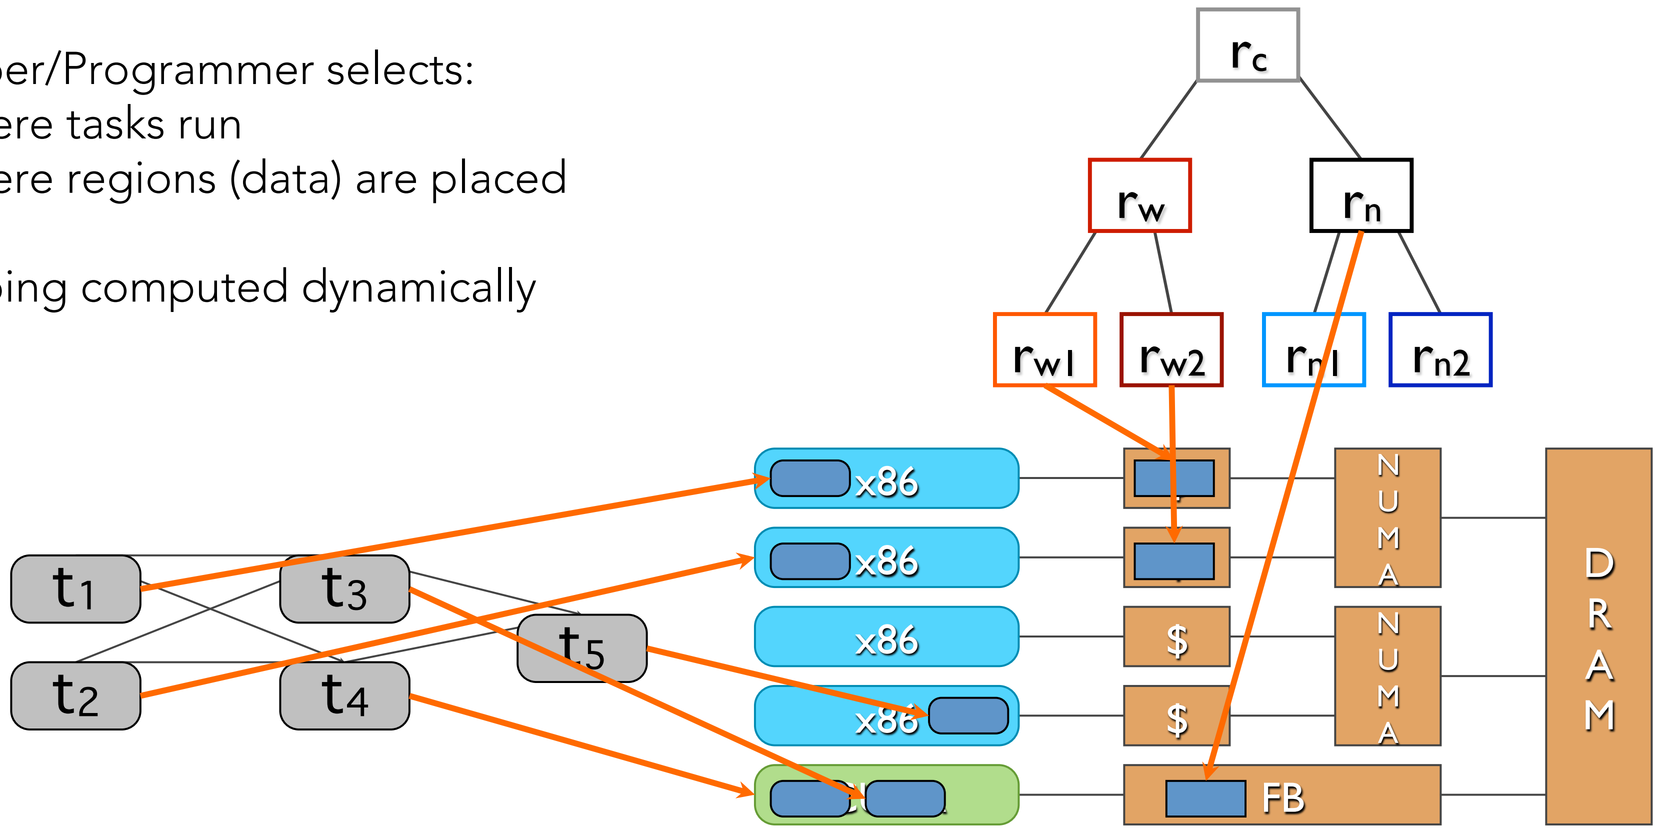

Legion on Titan

Titan node: 16-core AMD CPU
+ Nvidia K20X GPU

Sandia S3D Code in Legion

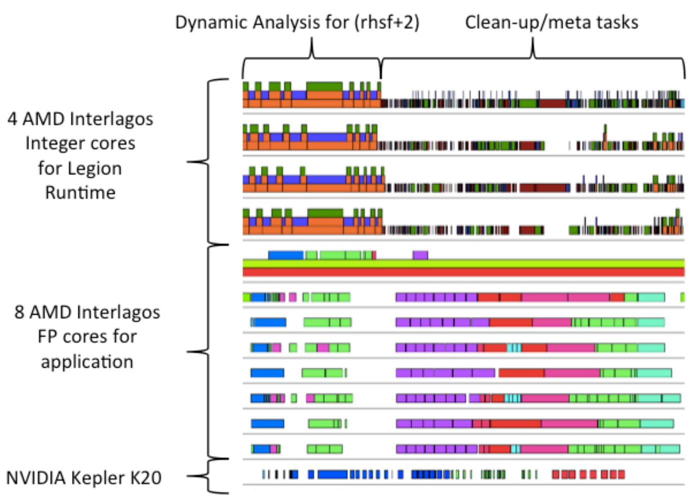

Mapping of Task/Data

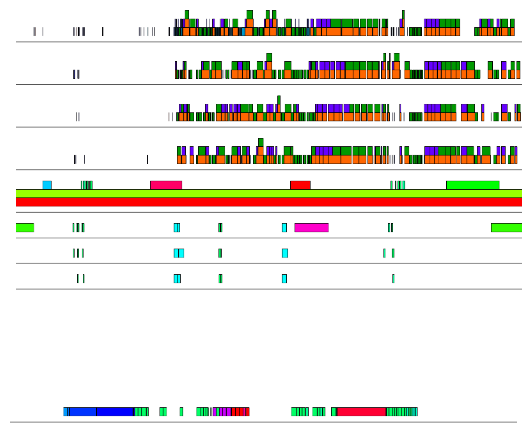

Option 2 - Simply a different mapper – code is the same

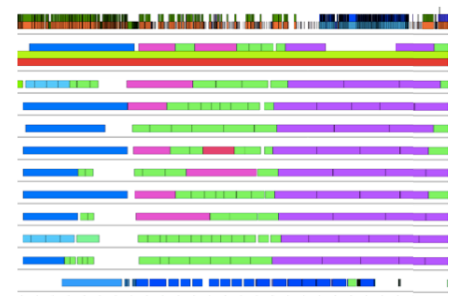

Finer grid – memory on GPU insufficient…

Stanford
PSAAP II

In conclusion…

In conclusion

Many scientific challenges…

- Physics: three-way coupling of turbulence/particles/radiation is largely unexplored; *is there a self-sustaining and robust cycle?*
- Software: each physics components has a favorite numerical method; *how to combine in a single, efficient and verifiable framework?*
- Exascale: DSLs have been demonstrated for single-physics problems; *can we combine multiple, interoperable DSLs and manage data efficiently and reliably on hybrid architectures?*
- UQ: aleatory and epistemic uncertainties are abundant and of different origin: *can we embed UQ techniques in the software infrastructure?*
- Validation: modern and detailed measurements are scarce; *can we design a controlled apparatus with large parameter variation to validate our tools?*

In conclusion

Lessons Learned:

- Three-way coupling (turbulence, particle dynamics, radiation) is critical

- …or is it four-way coupling: wall interactions and turbophoresis are critical

- …or is it five-way coupling: poly-dispersity is critical

- …or is it six-way coupling: CS+CS – integrated codesign team is critical

Dissemination

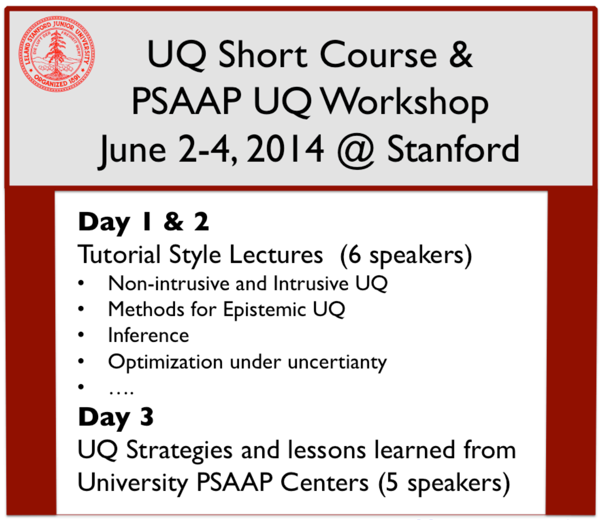

http://uq.stanford.edu

June 2-4, 2014 - ~100 attendees

- 16 from NNSA Labs
- 50 from Academia & Industry
- Stanford

Legion Bootcamp

http://legion.stanford.edu

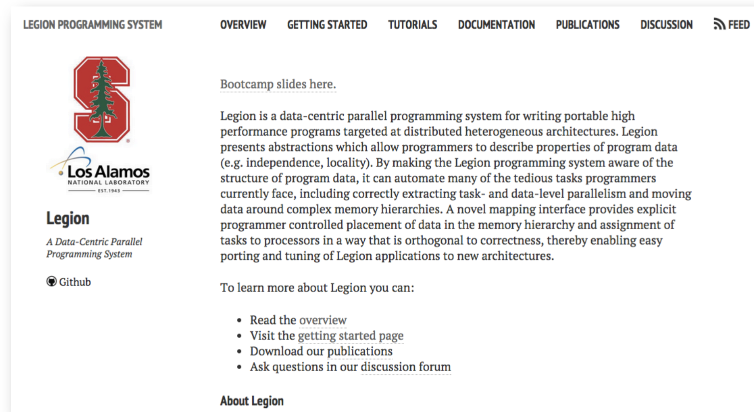

Dec 4, 2014 - 70 attendees

- Half from DoE (LLNL, Sandia, LANL, LBL, SLAC)
  Including a group from NNSA
- Industry (Intel, AMD, Nvidia)
- Stanford

Stanford PSAAP II

a collaboration between

Stanford University - University of Michigan – SUNY Stony Brook
University of Minnesota - University of Colorado - University of Texas

PSAAP II Project Plan

M1 Predictive simulations of low-Reynolds number turbulence with low mass loading under a moderate radiation input. Radiation is modeled using a lumped heat absorption model. Direct comparisons to Stanford's experiments will be carried out. Uncertainties account for the effect of imprecision in the specification of the boundary conditions and the parameters used in the absorption model. The current PSAAP codes, enhanced with Lagrangian tracking and the radiation model will be used.

M2 The objective in Year 2 is to increase the scope of the M1-type simulations by considering a more realistic uncertainty model for the absorption coefficient used in the radiation transport; a random field model will be considered. This will naturally require an increased number of computations; to enable these investigations, a DSL-enhanced version of the code will be deployed.

M3 In Year 3, the computational tools developed in the previous years will be exercised within a comprehensive UQ study. The objective is to investigate the effect of uncertainties in both the radiation source uniformity and particle size. In addition, the effect of the model uncertainty in the absorption model will be characterized and the resulting mixing efficiency variability will be characterized in detail. In this context the problem requires multiple M2-type simulations, further increasing the overall computational complexity.

M4 The objective in Year 4 is to perform simulations in a range of controlling parameters different from that considered in M1. Increases in the Reynolds number, the particle loading and a more realistic representation of finite-size effects in the particle transport will be considered. The use of resiliency libraries will be demonstrated.

M5 The goal in Year 5 is to revisit the UQ assessment carried out in Year 3 in the regime corresponding to the M4-type simulations. In addition to all the previous uncertainties, we will also consider additional sources related to the properties of the particles, for example emissivity imperfections. This will considerably enhance the computational complexity because of the increased dimensionality in the representation of the uncertainty. We anticipate to leverage the development in a probabilistic DSL to enable these final simulations.

More turbulence = Increase Re number = Eulerian grid size
More particles = Increase mass loading = Lagrangian tracking
More radiation = Increase heating = Strength of Coupling